xage SECURITY

Securing AI Systems with Identity-First Zero Trust

How enterprises can protect LLMs, agents, and AI workflows from emerging identity and access risks.

The New AI Security Crossroads

Enterprises are rapidly adopting GenAI, LLMs, and autonomous agents and as these systems connect across IT, OT, cloud, and edge, organizations arrive at a critical crossroads: AI is accelerating innovation, but also multiplying identity, access, and risks.

Exploding AI adoption across business functions creates new touchpoints for sensitive data.

Growing interdependence between users, agents, and systems means a single compromised identity can trigger cascading failures.

Connectivity protocols such as MCP and A2A make agents and models more capable, but also widen the attack surface when not paired with Zero Trust controls.

At this crossroads, enterprises must modernize their security strategy, shifting from user-only access models to verifying every identity, action, and interaction across the entire AI ecosystem.

IDC Insights Snapshot

What IDC Says About Enterprise AI Security Priorities

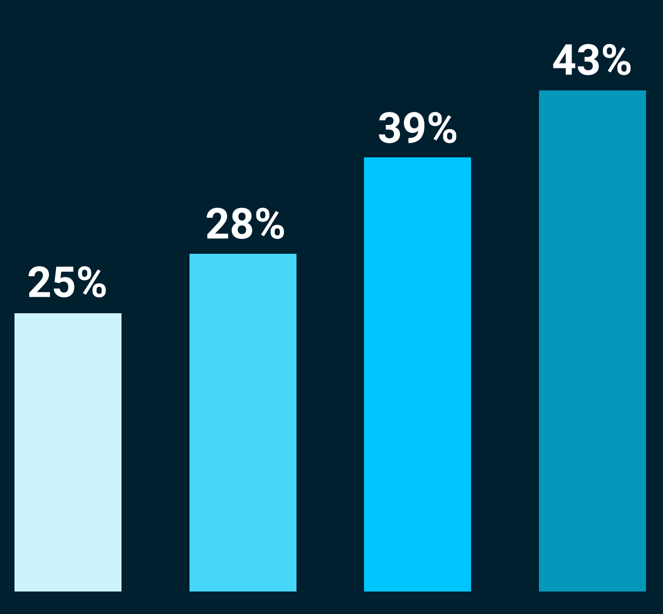

- **25%** cite authentication & authorization as top challenges.
- **28%** are expanding Zero Trust due to agentic AI.
- **39%** worry GenAI will cause data leakage.
- **43%** adopt Zero Trust for data security & governance.

IDC White Paper, sponsored by Xage Security, Securing AI with Zero Trust: Managing Identity & MCP Risks, #US53868625, October 2025.

Top Identity-First AI Risks

AI introduces new identity vulnerabilities that cannot be addressed with traditional access models

Identity Drift Across AI Pipelines
LLMs, tools, and agents accumulate overlapping entitlements that expand risk over time.

Agent Over-Permissioning
Autonomous agents often receive broad, persistent access that exceeds what is required.

Unrestricted Model-to-System Access
LLM functions and tools may interface directly with critical systems or sensitive data without strong authentication.

Chain-of-Agent Vulnerabilities
Actions triggered across multi-agent workflows can escalate privileges if not continuously verified.

Cross-Environment Credential Sprawl
AI components span IT, OT, cloud, and edge — each with different identity and access constraints.

Identity-First Zero Trust for AI

Enterprises need a Zero Trust approach that treats users, agents, and models as first-class identities. A modern AI Zero Trust architecture requires:

- **Unified Human + Machine Identity Control** across AI pipelines
- **Boundary enforcement** between LLMs, agents, services, and data stores
- **Dynamic, just-in-time entitlements** that eliminate standing privileges
- **Policy-driven agent behavior** with continuous verification
- **Context-aware guardrails** for sensitive prompts and system-initiated actions
- **Cross-domain enforcement** that spans IT, OT, cloud, and edge environments

Xage Zero Trust Fabric

Identity-first protection for AI, LLMs, and agents.

Unified Zero Trust
Secure every AI interaction.

Identity-First Controls
Every entity authenticated and authorized.

Jailbreak-Proof Enforcement
Blocks malicious prompts at protocol level.

Secure MCP & A2A
Least-privilege communication and tamperproof logging.

Dynamic, Just-in-Time Access
Eliminates standing privileges.

Integrated Governance
Redaction, audit trails, and NIST/EU AI Act compliance.

Resilient Fabric Overlay
Survives in disconnected, air-gapped environments.

Download the Full IDC White Paper

Learn how identity-first Zero Trust strengthens AI security across LLMs, autonomous agents, and enterprise workflows.

Securing AI with Zero Trust: Managing Identity & MCP Risks

IDC White Paper, sponsored by Xage Security, Securing AI with Zero Trust: Managing Identity & MCP Risks, #US53868625, October 2025.